Analysis of Technical Alternative Technologies for the Development of Context-Driven, Composable Environmental Representations for JSB

J. R. Hummel[*a], J. J. Bergenthal[b], W. F. Seng[c], J. R. Moulton[c,] and S.D. Prager[d],
[a]Argonne National Laboratory, Argonne, IL 60439
[b]Lockheed Martin Simulation, Training, and Support, 164 Middlesex Turnpike, Burlington, MA 01803
[c]JRM Technologies, 4500 Planck Road, Fredericksburg, VA 22407
[d]Lockheed Martin Simulation, Training, and Support, Bellevue, WA

ABSTRACT

The Joint Synthetic Battlespace for the Air Force (JSB-AF) is being developed to provide realistic representations of friendly and threat capabilities and the natural environmental conditions to support a variety of Department of Defense missions including training, mission rehearsal, decision support, acquisition, deployment, employment, operations, and the development of Courses of Action. This paper addresses three critical JSB issues associated with providing environmental representations to Modeling and Simulation (M&S) applications. First, how should the requirements for environmental functionality in a JSB-AF application be collected, analyzed, and used to define an Authoritative Environmental Representation (AER)? Second, how can JSB-AF AERs be generated? Third, once an AER has been generated, how should it be "served up" to the JSB-AF components? Our analyses of these issues will be presented from a general M&S perspective, with examples given from a JSB-AF centered view. In the context of this effort, the term "representations" is meant to incorporate both basic environmental "data" (e.g., temperature, pressure, slope, elevation, etc.) and "effects", properties that can be derived from these data using physics-based models or empirical relationship from the fundamental data (e.g., extinction coefficients, radiance, soil moisture strength, etc.) We present a state-of-the-art review of the existing processes and technologies that address these questions.

Keywords: Environmental representation, terrain materials, sensor development, modeling and simulation

1. INTRODUCTION

The environment can have a significant impact on military systems and operations. The images from Operation Iraqi Freedom of coalition forces moving through a driving sand storm and unable to see more than a few meters ahead provides a vivid example of how the environment can slow the pace of an operation and the performance of personnel and systems. When Modeling and Simulation (M&S) applications are developed to study the performance of systems and operations under "real world" conditions, it is imperative that **all** of the simulation entities, warfighting and environmental, be represented in as real a fashion that is consistent with the requirements and goals of the simulation.

If it is known that environmental factors are relevant in a given M&S application, a failure to **not** include an environmental representation or include one that is authoritative will mean that the simulation results will not be valid from a "real world" perspective and could have potentially significant consequences for the warfighter. When we refer to an environmental representation as being "authoritative", we mean that the processes, data, and technologies employed to generate it have been approved by subject matter experts (SME) and selected based on the context-specific requirements of the application.

1.1 The State of the Art in Developing and Providing Environmental Representations in M&S Applications

The overall requirements of a M&S application determine what set of environmental data and/or effects are required. The determination of these requirements has traditionally been performed as a manual, iterative process involving representatives from the M&S application or program office and SMEs from the environmental community. The process starts with a discussion of what the goals of the application are and what they mean from an environmental perspective. The iteration process between the M&S customer and the environmental community involves discussions of what the application would like to have, what it absolutely must have, what the environmental community can provide, and what,

[*]jhummel@anl.gov, phone 1 630 252-7189; fax 1 630 252-5128.

if any, constraints might be involved in the production process (e.g., cost, time, etc.). In performing this iterative process, it is critical that the end-user and the simulation developers be directly involved. In some previous efforts, the group with the responsibility to serve up the environmental data defined the environmental data requirements and the resulting environmental databases created were significantly larger than that actually required by the end-users.

The results from the process should include (1) an identification of the specific data elements that are required (e.g., temperature, pressure, winds, etc.); (2) their spatial and temporal resolution requirements; (3) the instructions on how and where the data should be delivered; and (4) the context of how they will be used in the application. Understanding the simulation context of how environmental data will be used is very critical because it provides insights to the environmental provider community of what kinds of sources they may have to deal with.

For example, current-art Air Force forward looking infrared and synthetic aperture radar sensors have enormous sensitivity, resolution, and fast update rates. Increasingly, imaging sensor simulators can render outputs at those high resolutions and update rates. The terrain elevation data and material classification detail must be at a resolution sufficient so as not to artificially limit the sensor in terms of background clutter. Here, the actual sensor resolution and update rates directly drive the AER spatial and temporal resolution requirements.

On the other hand, other visual M&S applications need only look realistic. In contrast to the previous example, the types of environmental data that must be provided would be vastly different. Unfortunately, the simulation context of how environmental data will be used is quite often not made available to the environmental community. Instead, environmental data providers are often just given a laundry list and told to provide everything they can, with the end result that the generated environmental databases are much larger than actually required by the simulation components.

Another reason the simulation context is important is that some environmental data elements can be produced from multiple sources and the context of what the parameter means and how it is intended to be used will govern which source will be used to generate the data. As an example, numerous applications require data on cloud cover amount and there are multiple sources for it. The context of how cloud cover will be used in an M&S application will provide critical insight into what approach should be used to generate the data. An aircraft entity that merely wants to know if it can see the ground only needs to know if there is a cloud between it and the ground. Where, specifically along that path the cloud exists, is not important. If on the other hand, the aircraft entity is doing a detailed target-background evaluation, detailed information on how much of the path is obscured, where the cloud is found along that path, and the density of water (or ice) within the cloud may be required.

Finally, programs ask for high spatial and temporal resolution databases knowing that the state of the art models makes the generation of these high resolution databases possible. However, an understanding of the simulation context may reveal that the high resolution data are not really necessary, or cannot be justified, because you may not be able to provide all of the required data at the same high level of resolution.

Within the Department of Defense (DOD) environmental community, three Modeling and Simulation Executive Agents (MSEA) have been chartered to address the issues associated with determining and meeting the environmental needs of the DOD community. (A fourth MSEA has also been charted for the intelligence community.) The three environmental MSEAs cover the air and space, ocean, and terrain domains. The terrain domain includes both natural terrain representations (e.g., soil, vegetation,…) and human created elements, like roads, buildings, etc. The environmental MSEAs provide guidance to DOD M&S programs on how to assess environmental requirements in their respective domains and provide mechanism and links to official operational DOD data providers that can supply the data (e.g., the Air Force Combat Climatology Center (AFCCC), the Navy's Fleet Numerical Meteorological and Oceanographic Center, or the National Imagery and Mapping Agency (NIMA)). Through the auspices of the Defense Modeling and Simulation Office's (DMSO) Integrated Natural Environment (INE) program, the MSEAs have developed a process called the Integrated Natural Environment Authoritative Representation Process (INEARP) that can be used to (1) analyze a customer's environmental requirements, (2) retrieve any archived representations if they exist and meet the requirements, (3) generate a customized AER if none is found in any archives, and (4) deliver the AER to the customer. The INEARP is expressed as a series of IDEF step-by-step processes that describe in great detail all of the steps required to go from a customer's statement of their requirements to an end product AER. (Details on the INE program and its concept of operations can be found in the document which is at the DMSO link: https://www.dmso.mil/public/library/projects/ine/inearpconops.pdf.)

1.2 Generating Authoritative Environmental Representations

Once the requirements have been established, AERs for two types of conditions have to be generated. The first type is a set of background environmental representations that can be generated in advance (e.g., basic terrain and atmospheric data.) The terrain data are typically predistributed and depending on the simulation requirements, the atmospheric data may also be able to be predistributed.

The second type of AER represents the environmental data or effects that need to be calculated during runtime as a result of dynamic interactions between simulation components (e.g., target-signature calculations, diffusion of battlefield obscurants, craters, etc.) The generation of these runtime representations will also require the use of authoritative environmental models or interactive databases.

The DOD environmental community has a number of operational sources for authoritative environmental data. Atmospheric and space data can be provided by operational Air Force and/or Navy agencies. For example, the Air Force Weather Agency and AFCCC are official providers of atmospheric data to the Air Force and their data products can be are available to support JSB-AF. For the terrain domain, data products from NIMA are used in the generation of terrain databases. In addition, these data products, which may include imagery products from a variety of sources, can be supplemented, as required, with data from civilian sources, such as the National Oceanic and Atmospheric Administration (NOAA) and the U.S. Geological Service (USGS).

As a part of the overall INE program effort, DMSO has funded the development of the Environmental Scenario Generator (ESG) program which was created to provide an architecture in which environmental requirements can be analyzed, environmental archives can be searched to see if any existing AERs meet a user's needs, and the generation of customized products if the archival searches yield no or insufficient results. ESG provides links to authoritative data and models in all four environmental domains. In the atmospheric domain, ESG provides links to authoritative DOD models, like MM5 and COAMPS (the operational mesoscale numerical weather prediction models used, respectively, by the Air Force and Navy) and the Advanced Climate Modeling Environmental System, a DMSO-developed system designed to provide historical weather data for any region of the world. In the space domain, ESG provides links to both DOD and civilian research grade and operational space models and data. In the terrain domain, ESG can provide links to NIMA data archives as well as the terrain generation tools that have also been developed by DOD. Finally, in the ocean domain, ESG can provide links to operational data and models maintained by the Navy (e.g., Naval Research Laboratory and Naval Oceanographic Office) and civilian agencies (e.g., NOAA). Finally, the resulting AERs found or generated by ESG can be delivered to the customer in a user-specified format, such as the Synthetic Environment Data Representation Interchange Specification (SEDRIS). (Details on SEDRIS can be found at https://www.dmso.mil/public/transition/sedris/.)

ESG includes a web-based interface that includes a set of sophisticated data mining tools that enable a customer to look for specific weather conditions that can be expressed in user-defined qualitative terms. For example, a user might want to have a data set that includes "hot and windy" weather conditions. With ESG's data mining tools, the user can define the precise quantitative conditions that denote "hot and windy" and ESG will search a global coarse resolution weather archive and identify weather data that match those conditions. The user can then choose to use these data as inputs into finer resolution numerical weather prediction models.

In the 1990's, the DOD operational environmental community was not always able to satisfy the M&S needs for AERs. As a result of significant DMSO sponsorship, the operational environmental community is now able to provide the requested AERs for DOD M&S applications and to provide them at <u>no cost</u> to official customers. The situation is now reversed in that more data can be supplied to DOD M&S customers than they can typically work with.

There are, however, at least two areas in which ESG may not be able to provide data necessary for JSB-AF applications. The first area involves a few atmospheric data elements required in radiative transfer calculations, like aerosol types, composition, and concentrations, that are not apart of the regular suite of parameters that are measured or modeled by the operational environmental community. However, the DOD and civilian meteorological communities have done extensive research in these areas that intelligent "guestimates" and inferences can be made to provide the required data if they are needed in JSB-AF applications. For example, researchers at the Air Force Research Laboratory at Hanscom AFB have developed a suite of very detailed aerosol models that were designed to support radiative transfer applications. These models cover all atmospheric altitudes from the surface to the upper atmosphere and a variety of phenomenological conditions like wind-lofted aerosols from different surface types[1,2], contributions from background stratospheric sulfuric acid[3] and volcanic eruptions, and inputs from meteoric dust.

The second area involves terrain material properties data that are required in a variety of areas, such as vehicle mobility and target-signature issues. Both the surface and sub-surface representations of these properties are critical in applications involving target-background signatures. At the most basic level, material properties which would be needed for the spectral bridge and its material radiance calculations would include thermo-physical properties, surface optical properties, and electromagnetic properties to cover the electro-optical, infra-red, and radio frequency (EO/IR/RF) bands. Examples of bulk properties affecting thermal band radiance include the density, specific heat, and thermal conductivity. Properties necessary for optical band radiance modeling would include solar absorptivity, emissivity, degree of Lambertian emission, polarization, as well as quantities directly used in modeling the reflectance distribution function such as the degree of specularly reflected light, the degree of diffusely reflected light and the diffuse hemispherical reflectance.

Other properties such as the specific coefficients used by a particular RF model would be necessary (for example, the coefficients in the Ulaby/Dobson radar cross section calculations) as dictated by the specific spectral bridge RF model. Due to transmissivity, the radiance of any particular material is governed not just by the material itself but also by what is underneath and proximate to it. The concept of material systems, comprised of layers of basic materials whose properties are enumerated in the database and boundary conditions between those layers, is essential to more closely modeling the material context which determines the RO/IR/RF band radiances of real-world terrain. Unfortunately, these data are difficult to obtain, especially the subsurface data, and their lack can negate the value of a high resolution terrain database. The "bottom line" is that just because you can produce high resolution terrain and weather databases, you may not always need them.

Current terrain database production technologies tend to focus on two primary areas of emphasis: accurate representation of terrain surface morphology and surface characteristics as they pertain to vehicle mobility. This focus is a legacy of the ground level heritage of many of the original M&S systems. As modern M&S systems continue to evolve, so do the applications of those systems to a variety of "new" problem areas. For example, a simulated behavior of a ground entity simply requires knowledge of the shape of the terrain surface and the manner in which that surface affects its mobility. On the other hand, a sensor's ability to detect that vehicle will be impacted by the thermal and emissive properties of the terrain surface and subsurface features surrounding the vehicle and the manner in which these properties vary as a function of the overall environmental conditions. There are extensive data on how some of these properties vary as a function of environmental conditions that have been collected by the general scientific community, but these data have typically not made it into the databases typically used in the DOD M&S community.

The training and doctrine community has successfully exploited a variety of probabilistic approaches in modeling the effects of rich environmental fidelity on sensor platforms. Though confusion matrices and other probabilistic approaches can be used with respect to the evaluation of doctrine, it is inappropriate to use such approaches when attempting to evaluate the performance of a sensor model itself. Consider the notion of a simulation based design cycle wherein a user wishes to test the relative efficacy of two differently polarized synthetic aperture radar systems for detecting particular phenomena in an urban environment. In this case, the characteristics of the environment itself serve as inputs to the model characterizing the sensor behavior. Without appropriately rich and realistic expression of the natural environment (i.e., "good enough for the task at hand"), the role of the environment in influencing system design and eventual acquisition cannot be adequately evaluated.

Undeniably, the material composition of both the terrain surface and anthropogenic features such as roads and buildings are directly relevant to a rigorous simulation-based design and acquisition cycle. It follows that efforts to better capture the material composition characteristics of the environment is required. Furthermore, depending on the system at hand, the opportunity exists to establish a trade space wherein fidelity in one domain (e.g., terrain surface morphology) can be traded for fidelity in another domain (e.g., materials composition) in an automated way and as appropriate to the requirements of the modeling effort.

1.3 Environmental Server Technologies

In past M&S applications, AERs were "served up" to the application components using two different approaches. In the first, the AER components were individually predistributed to the components. That is, a component that needed wind data got the wind components, while those that needed clouds got the cloud data, etc.

In the 1990's, when the Distributed Interactive Simulation (DIS) simulation approach, which is based on a set of Protocol Data Units (PDU), was being used to link distributed applications, an effort was begun to try and incorporate environmental effectslications. DIS is based on a set of. In the mid-1990's, a DIS Environmental Working Group was created to develop a set of environmental PDUs that could be used to distribute atmospheric and terrain data to DIS applications in support of the Synthetic Theater of War (STOW) program. These environmental PDUs were later incorporated into the Institute of Electrical and Electronic Engineers standard 1278.1a-1998 DIS protocols.

In support of STOW, a program was established called Weather in DIS (WINDS) to create a PDU-based environmental server that could be used to provide weather data to DIS applications. During the course of the STOW program, the HLA was developed and the WINDS system was modified to support the HLA by encapsulating the environmental PDUs as HLA object constructs. The WINDS system also evolved to be able to support environmental representation in the ocean and space domains and is now known as the Ocean Atmosphere and Space Environmental Server (OASES). OASES has been used to support a variety of DOD programs, such as the DMSO-sponsored series of three EnviroFed experiments and was also used in the initial JSB FY02 experiment.

Also during the 1990's, an object oriented simulated system called the Dynamic Environmental Effects Model (DEEM) was developed at the Argonne National Laboratory (ANL), a Department of Energy (DOE) Science Laboratory, to provide dynamic and integrated environmental effects to study the impact of the environment on operations. Initial DEEM applications involved a study for DOE on the impacts of different environmental remediation strategies at DOE "Superdump" sites and a study for the Joint Staff/J-8 of the impact on dynamic weather and terrain interactions on mobility, route planning, and the intelligent preparation of the battlefield. DEEM subsequently evolved into a government-owned and patented, subject domain independent framework called the Dynamic Information Architecture System (DIAS) that is being used in a variety of subject domains in both the military and civilian sectors. (When subject domain-specific models and applications are added to DIAS, it is typically renamed to represent the specific problem being address.) Table 1 provides a summary of a number of DIAS applications.

OASES and DIAS provide environmental representations in two totally different fashions. OASES provides data to federate components using the HLA publication and subscription services based on a Federation Object Model (FOM) that is fundamentally based on the original STOW environmental PDU designs. Object representations in point, 2-D, and 3-D forms are supported and they are typically updated at a specified (simulated) time interval. In a number of past studies that utilized OASES, such as the EnviroFed experiments, the atmosphere was represented as a full 3-D grid using data from the Navy's COAMPS model and runtime updates were provided at 20 (simulated) minute intervals.

In contrast, DIAS applications provide data and effects (in any subject domain) that are based on the specific simulation contexts of the individual components. There is no requirement for a common information model. Within DIAS, generalized domain objects work with a set of infrastructure objects in order to provide the specific representation (units, spatial form, etc.) that a simulation component requires.

In comparing OASES and DIAS, they are found to be fundamentally different systems that represent different ways in which environmental representations can be provided to M&S applications. OASES is basically a server that provides data, primarily via the HLA publication and subscription services. Using these services, OASES provides a data "push" mechanism. Information was found that indicated OASES can provide HLA interaction services (supporting a data "push" approach), but no examples have been found to demonstrate how they have been used in M&S applications.

The data served by OASES are typically generated externally. The resulting database, however, is often based on a common set of parameters and spatial and temporal formats rather than being constructed to meet the specific needs of a given application. For example, during the last EnviroFed experiment and the first JSB experiment, the OASES atmospheric database included the same set of parameters with the same temporal and spatial grid concepts – even though it was demonstrated that many of the requested parameters and spatial configurations were not required by any simulation component. Finally, the environmental object designs are still rooted in the original DIS PDU concepts which mean that users of the objects have the responsibility to transform the resulting data into any simulation specific formats.

Tab. 1. A Summary of Dynamic Information Architecture System Applications.

Dynamic Environmental Effects Model (DEEM) – DEEM was developed for the Joint Staff/J-8 and was used to perform dynamic weather and terrain influenced unit mobility, route planning and IPB analyses.
HealthSim – Developed for a private sector client, HealthSim provided an integrated physiological, clinical, logistical, and business oriented simulation of a health care delivery organization.
Distributed Integrated Ocean Prediction System (DIOPS) – Developed for the Naval Research Laboratory, DIOPS provides an integrated ocean architecture for a suite of a variable scale physics models that are used to support surf zone modeling for beach landing analyses.
Fort Future – DIAS is being used in the Fort Future effort to study issues associated with military base force protection, force training, base infrastructure development, force projection, and impact of base operations on surrounding areas.
Complex Adaptive System Countermeasures Analysis Dynamic Environment for Counter-Drug Applications (CASCADE-CD) – CASCADE-CD was developed for the Joint Staff/J-8 to provide an adaptive agent simulations for counter-drug operations analysis.
Virtual Ancient Mesopotamia – This DIAS application is part of an National Science Foundation-Grand Challenge to study agricultural and social sustainability of ancient urban centers.
CASCADE-FA - A CASCADE application of adaptive agent simulations for force application (FA) studies of joint operating concepts and architectures.

DIAS, on the other hand, is a simulation *framework* in which data and/or effects can be served and entities in any subject domain can dynamically interact with one another. DIAS is not tied to any specific data transmission protocol (e.g., DIS or HLA), but can support whatever is called for by the program's requirements.

DIAS domain objects are also designed from a generalized perspective and have demonstrated to be extensively reusable. Through the use of DIAS' infrastructure objects, entities can be provided with context-specific data and effects when and where it is important to them rather than waiting for updates to be published. DIAS can utilize the same authoritative provide data sources as OASES and can provide effects by integrating subject-matter approved models into its framework

In reviewing the potential environmental representation requirements of JSB-AF from the perspective of how the AERs should be provided to the applications, it is concluded that the requirements will vary significantly from one application to another in terms of (1) if data or effects (or both) will be required, (2) if data "push" or "pull" mechanisms are required, and (3) if the data need to be distributed using runtime protocols or if they can be predistributed. These requirements can be expressed simply as "One size will not fit all." In addition, discussions with numerous M&S customers have also indicated that they prefer to be able to interact with an environmental "black box" in order to obtain the data and/or effects where and when they need it – in the same manner that their real world systems do. It is further concluded that JSB-AF could be best served by an environmental architecture that is flexible and easily reconfigurable to be able to provide context specific environmental representations - an "Integrated Dynamic Environmental Architecture (IDEA)".

2. MYTHS AND REALITIES OF PROVIDING ENVIRONMENTAL REPRESENTATIONS

There have been a number of myths in the M&S community involving the inclusion of the environment in simulations[4]. The first myth is that including environmental representation is M&S applications is "too hard" because the environment is "so big." The reality is that the overall requirements of any simulation will determine the degree of difficulty required to incorporate any subject domain. Any subject domain entity may require a degree of fidelity that can tax a simulation's resources (e.g., fine spatial resolution, small temporal resolution, requirements for first principles physics, etc.) As a result, just because the playbox in which a simulation is taking place is "big" does not necessarily mean that the environmental representation required must also be "big."

A second myth is that including an environmental representation in M&S applications involves unique simulation requirements. The reality is that there are no unique simulation requirements for including environmental representations in M&S applications. Any subject domain entity can be created with a large number of elements requiring distribution among the simulation. The various data transmission protocol mechanisms (e.g., Distributed Interactive Simulation (DIS), High Level Architecture (HLA) Run Time Infrastructure, etc.) transmit "0's" and "1's" and do not know (or care) if a "1" or "0" is an environmental datum or not.

The final myth is that in the "real" world the environment is constantly changing and so any simulated environment must also be updated on a frequent basis. The reality is that the nature of how entities interact with the environment will dictate how often they need to know anything about the state of the atmosphere. An entity, like an aircraft, does not know (or even care) that the environment exists until in performs an action that involves the environment (e.g., sense, shoot, maneuver, land,…) In other cases, entities may only care if the environment evolves and crosses some specified thresholds that are important to it. For example, in a simulation that involves the launching and recovery of aircraft, changes in the state of the atmosphere may only become critical if they involve a change in flight rules. In a case like this, the information that could be required from the atmosphere is not a quantitative statement of what parameters have changed, but the qualitative impact of the change, such as a change from visual flight rules to instrument flight rules.

3. TECHNOLOGIES FOR INCORPORATING AERS IN M&S APPLICATIONS

OASES and DIAS provide environmental representations in two totally different fashions. OASES provides data to federate components using the HLA publication and subscription services based on a Federation Object Model (FOM) that is fundamentally based on the original STOW environmental PDU designs. Object representations in point, 2-D, and 3-D forms are supported and they are typically updated at a specified (simulated) time interval. In a number of past studies that utilized OASES, such as the EnviroFed experiments, the atmosphere was represented as a full 3-D grid using data from the Navy's COAMPS model and runtime updates were provided at 20 (simulated) minute intervals.

In contrast, DIAS applications provide data and effects (in any subject domain) that are based on the specific simulation contexts of the individual components. There is no requirement for a common information model. Within DIAS, generalized domain objects work with a set of infrastructure objects in order to provide the specific representation (units, spatial form, etc.) that a simulation component requires. This can be demonstrated by showing an example of how a DIAS domain object would be designed.

The atmosphere is one environmental domain that will be clearly required in JSB-AF. At the conceptual level, the atmosphere is a hydrodynamic fluid that can be represented with a variety of fundamental state parameters, like temperature, pressure, density, etc. These fundamental state parameters can also be used to derive a variety of additional parameters that are important to modelers, like relative or absolute humidity, cloud water content, pasquil stability, etc. Finally, additional atmospheric parameters, such as extinction coefficients, can be further derived to represent the effects of various atmospheric physics properties, like radiative transfer. The final set that is agreed upon will only define the set of states that can be provided to applications, but say nothing about how they will be provided (i.e., what set of units or spatial and temporal representations.) As will be discussed, these issues are addressed at runtime.

Once the total set of atmospheric state parameters has been decided upon, one must address the question of what kinds of atmospheric behaviors would be required to support JSB-AF applications. Examples of required atmospheric behaviors would be the ability to "evolve" (i.e., change over time), "attenuate radiation", and "transport" materials in the atmosphere, like smoke or battlefield obscurants. It is stressed that these are abstract statements of what the atmosphere can do and, again, provide no details of how it will be done in a given application (i.e., what specific model or algorithm). As with the state parameters, these details are addressed at runtime.

At runtime, DIAS infrastructure objects provide the runtime instantiations of object states and behavior. DIAS "Parameter" objects provide different point representations in variable formats while a set of "Spatial Data Set" (SDS) objects provide different 2-D and 3-D spatial representations, also in user-specified formats. Over twenty different SDS object concepts have been developed to cover a range of point and cell spatial representations. This is an important and innovative feature of DIAS and it addresses the critical reality that within a given application different entities may require the same atmosphere state parameter, but require it in a unique representation (e.g., units, spatial representation).

The specific mechanism by which an object's behavior is performed is also selected at runtime. This is because there can be different models or applications that could be used to represent an object's behavior depending upon the simulation conditions. The reality of this is demonstrated in a DIAS application called the Distributed Integrated Ocean Prediction System (DIOPS) that has been developed for the U.S. Navy.

DIOPS is being used to provide operational support for surf zone operations, namely the surf conditions that could be encountered during beach landings. These are calculated using a suite of authoritative models developed by the oceanographic community. Prior to the development of DIOPS, the suite of ocean models were manually run to calculate physical parameters in different ocean spatial domains and the results from one model were manually fed, via a classic "sneaker net", to the next model. Depending upon the conditions being considered, three different models could be used to provide the data on shallow water sea state dynamics. DIOPS was developed to automate the process and DIAS is used as the simulation framework in which all of the models reside. The context under which the various models can be used is known and captured within DIAS as a set of metadata that can be reasoned on to select the appropriate model for a given set of scenario conditions. With DIOPS, surf zone analyses can be performed in hours instead of days.

Models or applications that are used to represent object behaviors have been typically identified by program sponsors as the approved models and they are integrated into DIAS in an "as is" condition. That is, the models are implemented in whatever programming language they have been developed in and put into a software wrapper that provides the connections to and from DIAS. With the exception of providing software "hooks" into the models, there are no changes made to the internal aspects of the models. This "as is" integration philosophy successfully handles any proprietary issues associated with externally provided models. Custom built models can also be directly embedded into the DIAS infrastructure which can provide a more efficient level of connectivity.

The context of when the model can be used and how to use it is determined by analyzing a set of metadata that describes the input and output conditions and the ranges of applicability for the model. These metadata can be analyzed by DIAS in order to perform the runtime selection. Application programming interfaces (APIs) have been developed to enable users to add models on their own.

4. DEVELOPMENT OF AN INTEGRATED DYNAMIC ENVIRONMENTAL ARCHITECTURE

The technical rationale for developing the IDEA is to provide an environmental architecture that is representative of "real" world environmental concepts and responsive to the variable simulation needs of JSB-AF applications. In the "real" world, the environment is a highly dynamic and constantly evolving "black box". Entities exist in the black box, but are generally not "aware" that the environment exists until they require data from the environment or perform a behavior that involves an environmental effect. As an example, in a warfighter application an aircraft entity "flies" in the atmosphere, but the physics of flying are not simulated. The aircraft entity will not "know" that the environment is "there" until it performs an action that either requires atmospheric data or involves an atmospheric effect, like "sense", "shoot", or "land."

The IDEA would also be responsive to the needs of M&S developers who have indicated that their preferred method of interacting with the environment is to be able to "pull" the desired environmental data (or effect) out of the "black box" when, where, and how they need it as opposed to having it "pushed" to them at rates governed by the environmental server and in the structure of a consensus-developed information model. Finally, not all M&S applications require a runtime distribution of environmental data. In some past applications that involved the use of OASES, the length of the simulated scenarios was short enough that only one or two atmospheric data updates were published. In those cases, it could be argued that the environmental data could have been predistributed to the federates.

In developing the IDEA, the environment would be treated as a "black box" (Fig. 1) that could provide data and/or effects in a context-driven manner that is determined by the specific simulation needs of the entities that interact with the environment. The IDEA would be built using Argonne's DIAS as the architectural framework for supplying context-specific environmental representation to JSB-AF applications.

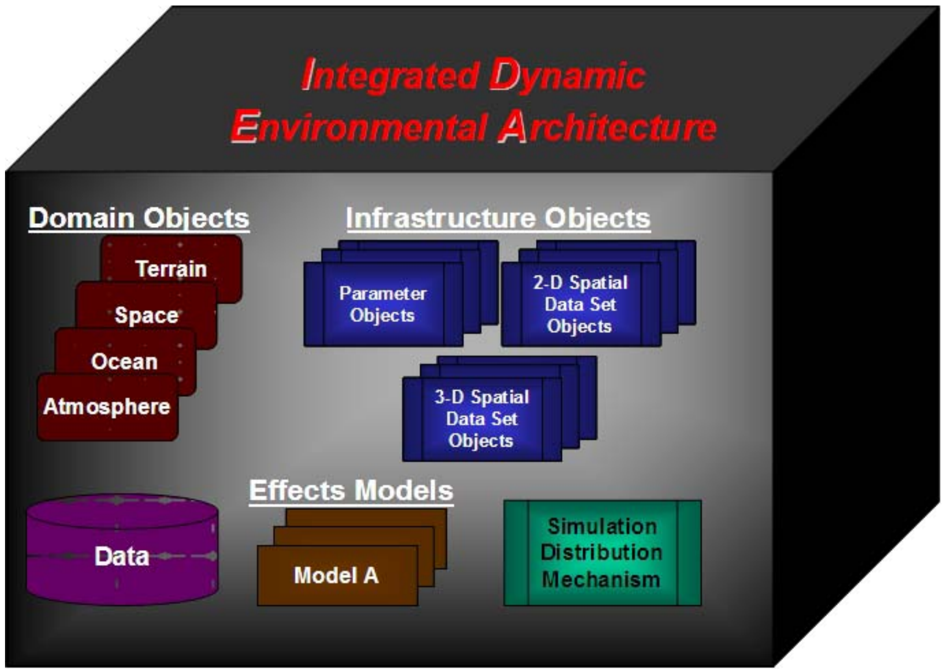

Fig. 1. A Schematic Representation of the Components of the Integrated Dynamic Environmental Architecture.

A set of generalized environmental domain objects would be provided that is rooted in the fundamental physics and concepts appropriate for each domain**.** Authoritative AERs supplied by ESG would be used to populate the domain objects. The DIAS infrastructure objects would be used to provide the runtime, context-specific instantiation of the objects.

When there is a requirement for environmental effects or data that must be calculated at runtime, SME-approved models would be integrated with the IDEA using the DIAS APIs. DIAS supports distributed operations so these models can be located on any networked platform.

We shall demonstrate how the IDEA can support the JSB-AF through a set of examples that represent potential JSB-AF applications. Figure 2 shows an example of the IDEA with a generalized atmosphere object consisting of a set of state parameters and three object behaviors. The atmospheric behavior of "*attenuateRadiation*" is shown being able to be represented by three models: MODTRAN[5], RADTRAN, and BACKSCAT[6]. These are authoritative models that have been developed by the Air Force Research Laboratory to address issues of atmospheric radiative transfer in visible and infrared, radar, and laser radar wavelengths, respectively. The conditions under which these models are valid are well understood and documented.

In the example, a sensor is querying the environmental black box to obtain a point radiance in the 8 – 12 μm wavelength from the sensor's location in space to a specified point A in space. From that query, IDEA can determine that the model MODTRAN will be required to perform the behavior of "*attenuateRadiation*" and that the atmospheric state parameters pressure, temperature, air density, and absolute humidity will be required along a line extending from the sensor's location to point A.

In the anticipated JSB-AF applications, we would envision that SME selected legacy models would continue to be the primary source of models of environmental effects. These models would be integrated into the IDEA using the DIAS APIs that have been developed to support model integration.

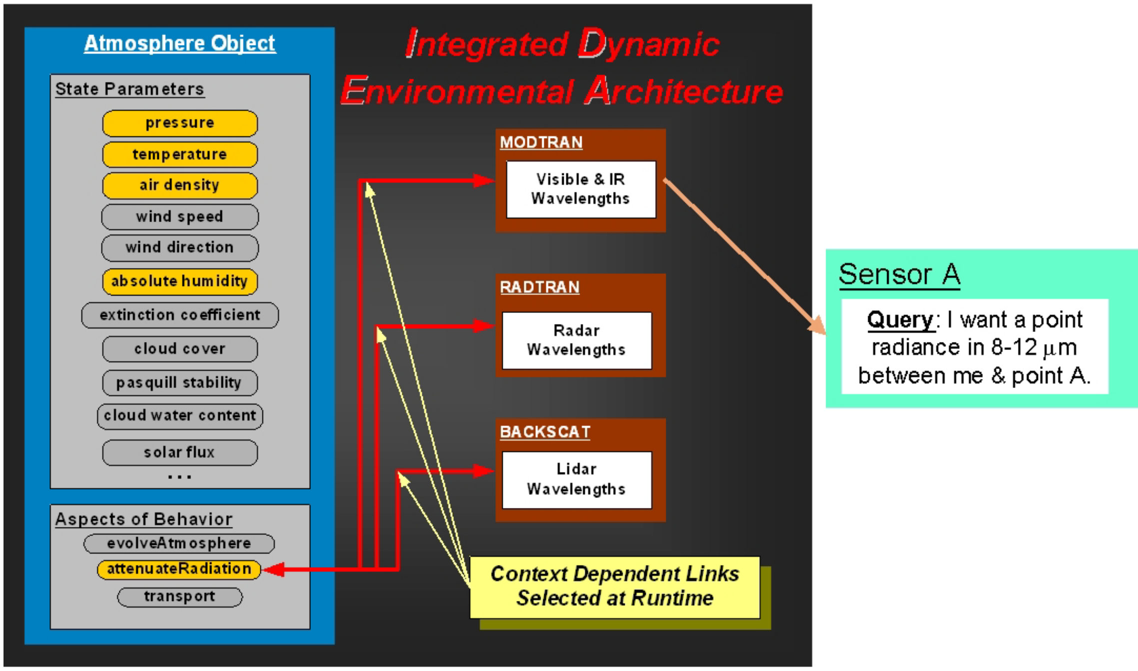

Fig. 2. An Example of how an Integrated Dynamic Environmental Architecture could Provide Context-Driven Environmental Representation to JSB-AF Applications.

5. RAPID GENERATION OF MATERIAL CODE ATTRIBUTED TERRAIN DATABASES

The organization and data structure of terrain databases which include material code assignments can be conceived of in several ways, but all begin with the user (or program) being able to assign a material code or codes to a given spatial area. Determination must first be made of how well the material at a given area must be known; that is its specificity: is it enough to know the material is broadly organic vs. metallic, or if organic, must the type of vegetation be known? Is a classification of generic trees or grass acceptable, or must the type of tree – e.g. conifer vs. deciduous - be known so that for example a physics-based approach to the calculation of at-aperture radiance can be performed based on the different optical and thermal properties of these two very different types of trees to generate a ray-traced scene? Once the specificity is known, the degree of accuracy and the ability to merge both with more specific and coarse databases must be determined.

Terrain database generation is essentially a fusion-based process wherein multiple data sources are utilized in a collective fashion to represent the physical environment. Inevitably, the final terrain database has some subset of the original information, reduced through either interpolation (in the case of the terrain surface) or through culling of attribution not specified in the requisite environmental data model (EDM). The EDM thus becomes a pivotal component in the development of materials coded terrain data, as it represents the "contract" between the producer and consumer of the environmental content.

There are multiple approaches with respect to representing the basis for a materials coded database in an EDM. In the simplest terms, objects characterizing a surface must contain attribution sufficient to embody their materials composition. This approach requires the enumeration of the material type (e.g., a roof object might be enumerated as "asphalt shingles"). The advantage of this approach is that it allows for a relatively straight forward database production process, in that the materials coding schema is tied to a finite set of materials types. Importantly, by itself, this level of representation is inadequate to meet the needs of modern high fidelity sensor platforms as these platforms have very specific demands in terms of their capabilities relative to their bandwidth and secondary properties of the object (dielectric, specular and thermal characteristics).

It is possible to represent higher fidelity information in the EDM (e.g., the spectral reflectance of an object in a given bandwidth). Unfortunately, such an approach requires large amounts of supporting data - as inputs to the terrain database generation process, these data are not often (or easily) available. For this reason, a system designed to rapidly construct terrain databases in a way that is capable of capturing materials coding is required. The materials coding, in turn, supports simulation based design and acquisition, but not without a mechanism to facilitate the bridge between the object with an enumerated type and the more robustly defined spectral information needed by the sensor model. This spectral bridge serves to fuse the information required for sensor processing both when and in the manner required.

The spectral bridge concept is generally well aligned with the desired "pull" architecture envisioned for the various JSB programs. Essentially, the bridge mechanism serves as means by which any given object and its materially coded surfaces are arbitrated against the required environmental objects for ambient conditions and a set of spectral objects containing full spectrum emissivity, transmissivity and other bandwidth related models. As the spectral objects inherit the environmental conditions and the surface materials information, they then become the focal point for sensor/environment interaction.

The advantage of the spectral bridge is that the only data requirements are standardized enumerations for surface materials types. Though the current SEDRIS Environmental Data Coding Specification (EDCS 3.1) has sufficient representational capacity to characterize the bandwidth specific property sets (e.g., RADAR_MATERIAL_PROPERTY_ SET, INFRARED_MATERIAL_PROPERTY_SET and others), arbitrarily attaching these data to a terrain database serves to enlarge a terrain database with data specific to a very limited number of users. Segmenting the atmospheric, terrain and spectral data thus leads to a more partitioned problem space with less significant interdependencies should any related part of the overall architecture be subject to change.

Assuming a database construction paradigm wherein materials coded attributes are a primary requirement leads to consideration of the origin of such information. Generally, materials information are collected from an assortment of multi-spectral platforms such as satellite based broadband sensor systems, hyper-spectral platforms, radar based surface roughness analysis and direct observation. These data are generally represented as continuous, though categorically partitioned surfaces. Importantly, data of this sort tend to characterize surface materials and not necessarily materials composition. As such, if materials composition is required (e.g., substrate information), these data must be derived either heuristically, or through inference based on the combination of surface type and other available factors. Regardless of how surface and substrate information are collected/derived, they are very dense data with an *n*-squared increase in size relative to resolution.

Rapid generation of material coded terrain databases needs to be considered from both the terrain perspective and the anthropogenic feature perspective. From the terrain surface perspective, the background material characteristics are generally used to aid in the determination of some type of state (e.g., detection of a stationary vehicle, differentiation of concealment from natural phenomena) or some type of state change (e.g., the detection of a moving vehicle). The important concern becomes the utilization of the materials information most suited to the desired analytical process.

Typically cumbersome database generation processes coupled with variable needs suggests that a methodology capable of integrating materials information in something of a near real-time basis would be an appropriate strategy. Notionally, this aligns with the objective nature of the envisioned systems (again, a pull oriented system). Furthermore, it allows potentially disparate data sets to be used in conjunction with one another. A low resolution terrain database could be used with high resolution materials information without having to compromise the overall performance gained by maintaining a lower resolution terrain surface. Finally, this approach facilitates the utilization of nearly any information characterizing terrain surface materials, assuming the materials information can be mapped to the identified set of standard enumerations.

Anthropogenic objects, specifically buildings, are also subject to significant sensor scrutiny. In that the sensors used to examine buildings and similar objects include IR, radar, fluoroscopic and acoustic platforms, both surface and substrate material information becomes extremely important. Unlike terrain, the compositional characteristics of the object are directly associated with the object itself. Additionally, whereas each "piece" of terrain is essentially a single faceted surface, features such as building walls may have differing exterior, interior and internal compositions. Unfortunately, there is a significant shortage of accessible information characterizing individual structures.

With regard to terrain surfaces, the emphasis is on the handling and enhanced integration of available materials information into an otherwise difficult database generation process. This approach facilitates a "mix and match" approach to the integration and construction of simulation ready natural environments, thereby allowing a software capability to handle the relationship between the materials information and the actual terrain database. With regard to objects such as buildings, the proposal is to evaluate the use of a multifaceted attribute derivation process to "fill in the gap" where required attribution is not available. This approach enables specific treatment of individual objects, coherent treatment of groups of objects, and is based on a proven process that can be tailored to JSB-AF needs.

JRM Technologies' internal and government-sponsored research has led to the design and implementation of a first order material classifier. The classifier examines the spectral nature of pixels in a scene (hyperspectral, red-green-blue (RGB), grayscale, etc.), and assigns a material property to each pixel without human-in-the-loop interaction. The methodology utilized by the classifier is that the at-aperture radiances of a candidate list of materials is calculated by the Signature

Predictive Classifier (SPC) utilizing physics principles and the materials' fundamental physical properties, and the algorithm then finds the best match between the radiance in the pixel in the scene and the radiance of a material in the list. Current implementation of this allows the assignment of one material per pixel, and development is ongoing to extend this to an arbitrary n-material assignments per pixel. The more spectral information there is in the scene to be classified, the more rigorous and exact the material classification of pixels in the scene.

The material classifier concept can be extended to additionally include a rule set which can help to overcome the limitations imposed by the lack of data that is inherent in any picture. This rule set, referred to as a Spatial Context Refiner (SCR), looks at the neighboring pixels' material classification to insure that, for example, darkened blotches of shadowed grass in a field do not become classified mistakenly as asphalt sprinkled in the field. Lastly, the Iterative Context Adjuster (ICA) will be designed to allow iteration of the governing constraints applied to the SPC and SCR. With the user guiding the extent of allowable constraint variation, the ICA will iteratively re-call the SPC and SCR to determine the highest confidence prediction. Of particular note for thermal-band data, the ICA will also allow variations to be explored in the material systems library. This can be useful for thermal images where the thickness of individual materials in the material system are not precisely known.

One of the reasons that the SCA and ICA are necessary even after classification is that material system assignment may not match with ground-truth exactly, even if the physical parameters are perfectly specified, since many effects, both intrinsic to and outside of the material system, can affect the response. Natural variation in the spectral response of a single material can cause it to resemble another. For example self-shadowing within a forest canopy can cause tree leaves to be misclassified as asphalt. Similarly, natural variation can cause two different material systems to resemble the same material system. For example, blue paint can resemble water reflecting the sky. This is a significant problem in images with a low number of channels. Increasing the number of channels reduces this problem, but does not eliminate it because the variation can be large enough that the error can even fool even high data-density systems.

Lockheed Martin Information Systems Advanced Simulation Center has also executed substantial internal research and development leading to the design and implementation of a series of tools enabling the heuristic and inferential generation of terrain and feature attribution. The process devised facilitates attribution of objects based on a combination of rule sets constructed for the area of interest, information such as the Urban Terrain Zones, available attribution and other geotypical information. The resultant output is a series of objects (the methodology is generally exercised on buildings, but is equally applicable to other feature types as well) with fully populated attribution that is geotypical and appropriate to the area of interest.

Once the terrain database (objects or terrain surface) is populated with attribution meeting the schema defined in the requisite environmental data model, it is suitable for use in conjunction with the spectral bridge. The spectral bridge is essentially a software interface that fuses available spectral data with terrain database content and its associated materials coding attribution. Whereas a conventional synthetic natural environment requires that the spectrum-specific sensor information and the terrain and/or object information reside in a single instantiation of the environmental database, the bridged spectral and environmental information may be distributed (thereby also resolving potential multilevel security issues). Furthermore, as the spectral bridge has knowledge of the terrain, spectral characteristics, and sensor (e.g., resolution, emissivity and sensor bandwidth, respectively), it is capable of optimizing the returned signatures relative to the requirements of the composed system.

In its elemental form, the spectral bridge is an interface that requires a series of specific inputs: the attributed terrain surface or objects, an environmental data model, spectral data, sensor bandwidth, and any sensor or regionally specific models. Collectively, these elements define the basis for rapid integration of materials coded environments into the JSB architecture as summarized in Table 2.

Figure 3 gives a schematic representation of the processes involved in developing Material Coded Attributed Terrain Databases (MCAT-DB). In examining the processes, it is noted that the majority of the processes are performed pre-runtime. The two that are performed at runtime are those related to any interactions with sensor specific models or any other simulation specific processes, like the dynamic impact of environmental processes (e.g., changes in soil properties as a result of precipitation or evaporation.)

The final analysis shall include an example terrain database, multiple examples of terrain materials information at varying resolutions, the spectral bridge, an example set of attributed buildings and an example library of spectral information. Importantly, the study will include an analysis of the JSB representational requirements and shall capture these requirements in an environmental data modeling tool, such as the Common Data Model Framework.

Tab. 1. A summary of the elements and roles they play in developing material coded environmental representations.

Element	Role
Attributed surface/objects	The basic unit of information extracted from the synthetic natural environment and input into the spectral bridge. Contains enumerated materials codes, either algorithmically derived or populated from authoritative sources.
Environmental Data Model	The defined set of features, attributes and allowable values. The set of enumerates characterizing materials codes are defined in the EDM.
Spectral data & mappings	The core of the spectral bridge capability. A set of spectral information aligned with the EDM defined enumerates for materials codes. These data may vary regionally (e.g., asphalt shingles may have different composition in CONUS v. OCONUS locations). The spectral data represents an evolving library that enables more universal application of spectral information (as opposed to binding it to a single instance of a database).
Sensor bandwidth	The portion of the electromagnetic spectrum that requires consideration during any given instantiation of the bridge. Used to optimize the query, processing and return of spectral characteristics.
Secondary models	Facilitates an extensible system wherein external factors (e.g., environmental conditions, platform specific capabilities) can be used to adjust the reported data. This capability is important to the simulation based design and acquisition cycle as it allows ready alteration of state parameters (either static or dynamic), enabling a variety of bootstrap analysis paradigms. The secondary models also enable the return of complex signatures wherein multiple materials codes must be considered (e.g., walls, camouflage, etc.).

6. A PROPOSED SOLUTION

Figure 4 provides an example of how the IDEA and the MCAT-DB and tools could be used to provide target-signature information to an aircraft that is trying to target a tank that is on a basic terrain background. In this example, four objects are being represented: an *Aircraft*, *Atmosphere*, *EarthSurface*, and a *GroundVehicle* object. Each conceptualized object has, at least, one or more behaviors that they can perform. The behaviors are linked to a model or process that can provide the runtime instantiation of the behavior. (In this example, we have deliberately not used a "brand name" identifier for what models are being used nor have we detailed how the information goes from one object to another (i.e., what data transmission protocol is used.) In both cases, these are details that would be determined as a part of the overall requirements definition phase.)

The box in the lower right hand corner is where the MCAT-DB and tools would reside. The database icon labeled "Material Properties Database" represents the base set of material properties (e.g., emissivities, conductivities, etc.) that would be defined in terms of a set of default environmental conditions (e.g., high noon, clear skies, dry soil conditions, etc.) This default set of data would be used with a set of tools that are included in the model box called "Material Properties Mapper" that would take the default data and specific environmental conditions and transform them into the environmentally context specific properties required in the target-background signature calculations.

7. ACKNOWLEDGMENTS

This work is sponsored by the U.S. Department of Energy under contract W-31-109-ENG-38.

REFERENCES

1. Shettle, E.P., and Fenn, R.W. (1979) Models of the Aerosols of the Lower Atmosphere and the Effectgs of Humidity Variations on Their Optical Properties, Air Force Geophysics Laboratory, Hanscom AFB, MA, AFGL-TR-79-0214.
2. Longtin, D.R., Shettle, E.P., Hummel, J.R., and Pryce, J.D. (1988) A Wind Dependent Desert Aerosol Model: Radiative Properties, Air Force Geophysics Laboratory, Hanscom AFB, MA, AFGL-TR-88-0112.
3. Hummel, J.R., Shettle, E.P., and Longtin, D.R., (1988), A New Background Stratospheric Aerosol Model for Use in Atmospheric Radiation Models, Air Force Geophysics Laboratory, Hanscom AFB, MA, AFGL-TR-88-0166.
4. Hummel, J.R. (1996) "Addressing Questions About Including Environmental Effects in the DMSO HLA," Proceedings of the 15th Workshop on Standards for DIS, Institute for Simulation and Training, Orlando, FL, September 1996.
5. A. Berk, G.P. Anderson, P.K. Acharya, J.H. Chetwynd, M.L. Hoke, L.S. Bernstein, E.P. Shettle, M.W. Matthew, and S.M. Adler-Golden, MODTRA4 Version 2 User's Manual, Air Force Research Laboratory, Hanscom AFB, MA, 1 June 1999 (Revised 17 April 2000).
6. J.R. Hummel, D.R. Longtin, N.L. DePiero, and R.J. Grasso, "BACKSCAT Lidar Simulation Version 3.0: Technical Documentation and Users Guide," Phillips Laboratory, Hanscom AFB, MA, PL-TR-92-2328, 3 December 1992.

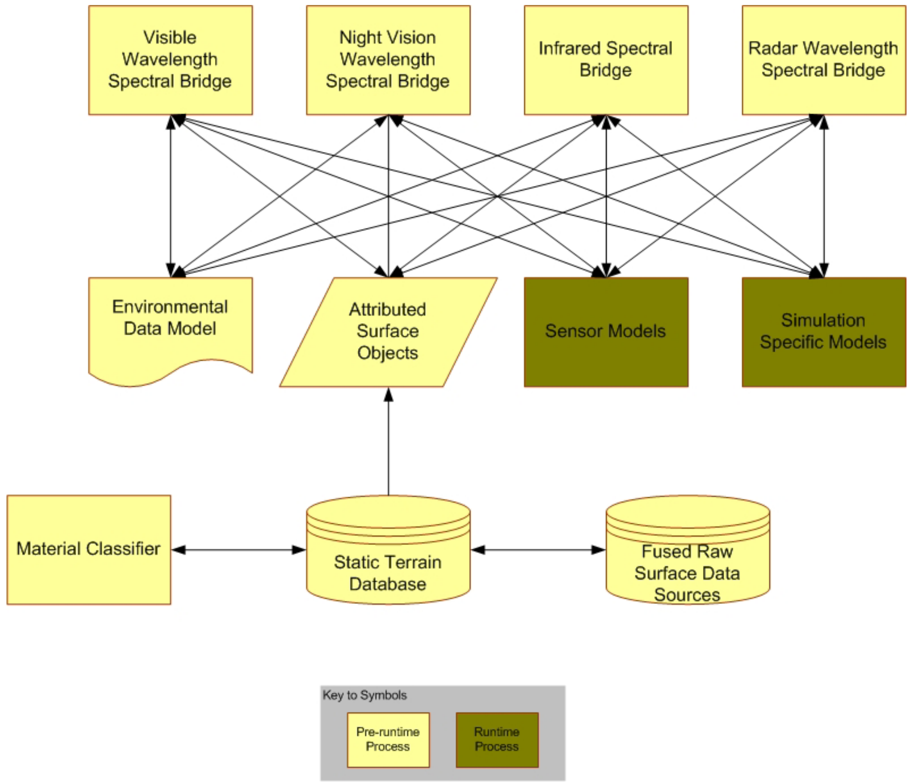

Fig. 3. Schematic representation of the processes involved in the generation of material coded attributed terrain databases.

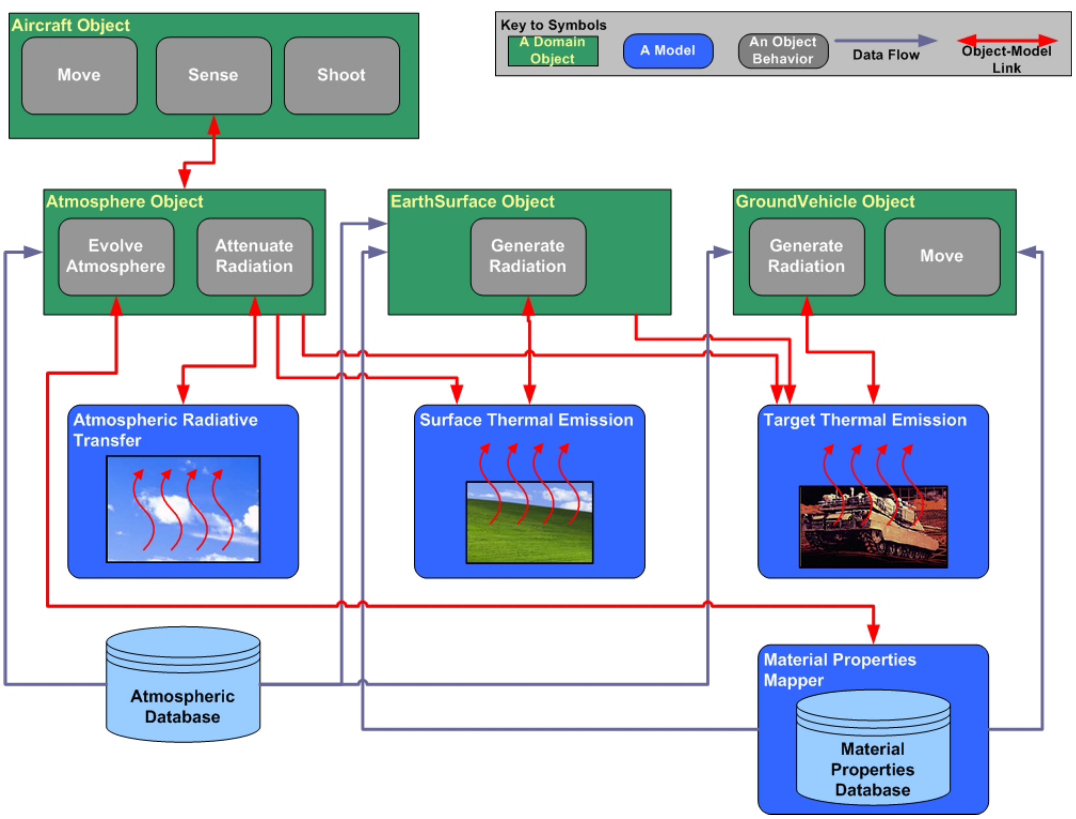

Fig. 4. A schematic representation of how the IDEA and the material coded attributed terrain database and tools would be used to provide target-signature data to an interacting aircraft object.